Universal Joint Image Clustering and Registration using Partition Information

Ravi Kiran Raman and Lav R. Varshney

arXiv:1701.02776v1 [cs.IT] 10 Jan 2017

Abstract

The problem of joint clustering and registration of images is studied in a universal setting. We define universal joint clustering and registration algorithms using multivariate information functionals. We first study the problem of registering two images using maximum mutual information and prove its asymptotic optimality. We then show the shortcomings of pairwise registration in multi-image registration, and design an asymptotically optimal algorithm based on multiinformation. Finally, we define a novel multivariate information functional to perform joint clustering and registration of images, and prove consistency of the algorithm.

Index Terms

Image registration, clustering, universal information theory, unsupervised learning, asymptotic optimality

I. Introduction

Suppose you have an unlabeled repository of medical images of MRI, CT, and PET scans of brain regions corresponding to different patients from different stages of the diagnostic process. You wish to sort the images into clusters corresponding to individual patients and align the images within each cluster. In this work, we address this exact problem of joint clustering and registration of images using novel multivariate information functionals.

Different digital images of the same scene can appear significantly different from each other based on the imaging device, the application of filters, and orientation of the images. Such meta-data about the digital images is often not completely available. This emphasizes the need for universality in the design of reliable clustering and registration algorithms.

The tasks of image clustering and registration have often been dealt with independently. We emphasize here that the two problems are not necessarily independent and define a universal, reliable, joint clustering and registration algorithm.

There is a rich literature on clustering and registration; we describe a non-exhaustive listing of relevant prior work.

Unsupervised clustering of objects has been studied under a wide variety of optimality and similarity criteria. The k-means algorithm and its generalization to Bregman divergences [1] are some popular distance-based methods..

In this work, we focus on information-based clustering algorithms [2], [3], owing to the ubiquitous nature of information functionals in universal information processing. The task of universal clustering has been studied in the communication and crowdsourcing settings in [4] and [5], respectively.

Separate from clustering, multi-image registration has been studied extensively. Some prominent region-based registration methods include maximum likelihood (ML) [6], minimum KL divergence [7], correlation detection [8], and maximum mutual information (MMI) [9]. Feature-based techniques such as [10], [11] have also been explored. For a more comprehensive and detailed survey of existing techniques, see [12].

Lower bounds on mean squared error for image registration in the presence of additive noise using Ziv-Zakai and Cramer-Rao bounds have been explored recently [13], [14]. The MMI decoder was originally developed in communications [15] and deterministic reasons for its effectiveness in image registration have also been studied [16]. Correctness has also been established through information-theoretic arguments [17].

While the MMI method has been found to perform well through numerous empirical studies, concrete theoretical guarantees of optimality are still lacking. In this work, we extend the framework of universal delay estimation for memoryless sources [18] to derive universal asymptotic optimality guarantees for the MMI method under the

This work was supported in part by Air Force STTR Contract FA8650-16-M-1819.

The authors are with the University of Illinois at Urbana-Champaign.

Hamming loss. Further, we show the shortcomings of pairwise information measures in multi-image registration and define novel multivariate functionals that overcome them.

We next introduce the image and channel model, and characterize rigid-body transformations. We then prove asymptotic optimality of the MMI method for two-image registration. We then consider multi-image registration and define asymptotically optimal registration algorithms using multiinformation. Finally, we perform reliable, universal, joint clustering and registration of images using multivariate information.

II. Model

We now formulate the joint image clustering and registration problem and define the model of the images that we work with.

A. Image and Noise

In this work, we consider a simple model of images. Specifically, we assume that each image is a collection of n pixels drawn independently and identically from an unknown prior distribution $P_R(\cdot)$ defined on finite space of pixel values $[r]$. Since the pixels are drawn i.i.d., we represent the original scene of the image by an n-dimensional random vector, $\mathbf{R} \in [r]^n$. More specifically, the scene is drawn according to $\mathbb{P}[\mathbf{R}] = P_R^{\otimes n}(\mathbf{R})$.

Consider a finite collection of ℓ distinct scenes (drawn i.i.d. according to the prior) $\mathcal{R} = \left\{\mathbf{R}^{(1)}, \ldots, \mathbf{R}^{(\ell)}\right\}$. This set can be interpreted as a collection of different scenes. Each image may be viewed as a noisy depiction of an underlying scene.

Consider a collection, $\{\tilde{\mathbf{R}}^{(1)}, \ldots, \tilde{\mathbf{R}}^{(m)}\}$, of m scenes drawn from $\mathcal{R}$, with each scene chosen independently and identically according to the pmf $(p_1, \ldots, p_\ell)$. Specifically the image i corresponds to a depiction of the scene $\tilde{\mathbf{R}}^{(i)}$, for $i \in [m]$.

We model the images corresponding to this collection of scenes as noisy versions of the underlying scenes drawn as follows

$$\mathbb{P}\left[\tilde{\mathbf{X}}^{[m]} \middle| \tilde{\mathbf{R}}^{[m]}\right] = \prod_{j=1}^{\ell} \prod_{i=1}^{n} W\left(\tilde{X}_i^{(K(j))} | R_i^{(j)}\right), \tag{1}$$

where $K(j) \subseteq [m]$ is the inclusion-wise maximal subset such that $\tilde{\mathbf{R}}^{(i)} = \mathbf{R}^{(j)}$ for all $i \in K(j)$. That is, images corresponding to the same scene are jointly corrupted by a discrete memoryless channel, while the images corresponding to different scenes are independent conditioned on the scene. Here we assume $\tilde{\mathbf{X}} \in [r]^n$. The system is depicted in Fig. 1.

B. Image Transformations

The corrupted images are also subject to independent rigid-body transformations such as rotation and translation. Since images are vectors of length n, transformations are represented by permutations of $[n]$. Let Π be the set of all allowable transformations.

Let $\pi_j \sim \text{Unif}(\Pi)$ be the transformation of image j. Then, the final image is $\mathbf{X}_i^{(j)} = \tilde{\mathbf{X}}_{\pi_j(i)}^{(j)}$, for all $i \in [n]$. Image $\mathbf{X}$ transformed by π is depicted interchangeably as $\pi(\mathbf{X}) = \mathbf{X}_\pi$.

In order to facilitate registration, we assume that we are aware of Π. Further, we assume Π forms a *commutative algebra* over the composition operator $\circ$. More specifically,

1) for $\pi_1, \pi_2 \in \Pi$, $\pi_1 \circ \pi_2 = \pi_2 \circ \pi_1 \in \Pi$;
2) there exists unique $\pi_0 \in \Pi$ s.t. $\pi_0(i) = i$, for all $i \in [n]$;
3) for any $\pi \in \Pi$, there exists a unique inverse $\pi^{-1} \in \Pi$, s.t. $\pi^{-1} \circ \pi = \pi \circ \pi^{-1} = \pi_0$.

The number of distinct rigid transformations of images with n pixels on the $\mathbb{Z}$-lattice is polynomial in n, i.e., $|\Pi| = O(n^\alpha)$ for some $\alpha \leq 5$ [19].

Definition 1 (Permutation Cycles): A *permutation cycle*, $\{i_1, \ldots, i_k\}$, is a subset of permutation π of $[n]$, such that $\pi(i_j) = i_{j+1}$, for all $j < k$ and $\pi(i_k) = i_1$.

It is clear from the pigeonhole principle that any permutation is composed of at least one permutation cycle. Let the number of permutation cycles of a permutation π be κ_π.

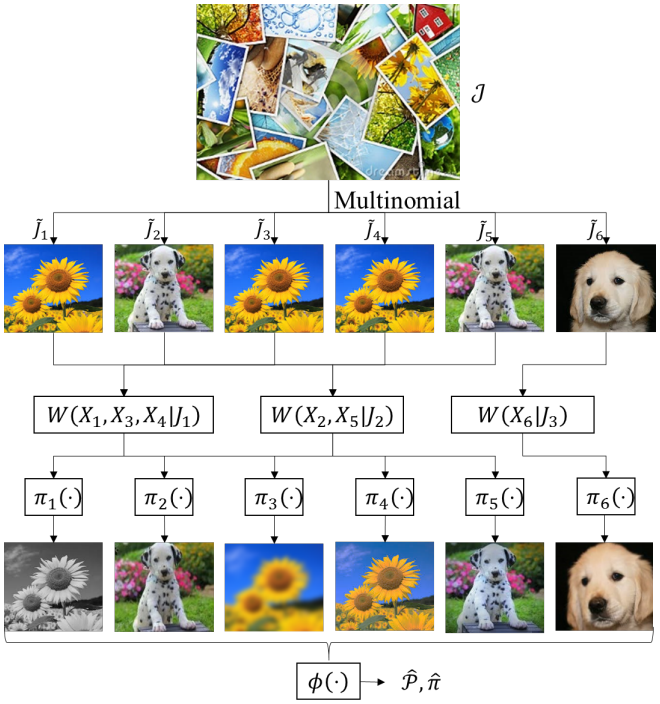

Fig. 1. Model of the joint image clustering and registration problem

Definition 2 (Identity Block): The *identity block* for any permutation $\pi \in \Pi$ is the inclusion-wise maximal subset $\mathcal{I}_\pi \subseteq [n]$ such that $\pi(i) = i$, for all $i \in \mathcal{I}_\pi$.

Definition 3 (Simple Permutation): A permutation π is a *simple permutation* if $\kappa_\pi = 1$ and $\mathcal{I}_\pi = \emptyset$.

Definition 4 (Non-overlapping Permutations): Any two permutations $\pi, \pi' \in \Pi$ are said to be *non-overlapping* if $\pi(i) \neq \pi'(i)$ for all $i \in [n]$.

Lemma 1: Let π be chosen uniformly at random from the set of all permutations of $[n]$. Then, for any constants $c \in (0, 1]$, C,

$$\mathbb{P}\left[|\mathcal{I}_\pi| > cn\right] \lesssim \exp\left(-cn\right), \quad \mathbb{P}\left[\kappa_\pi > C\frac{n}{\log n}\right] = o(1). \tag{2}$$

Proof: First, we observe that the number of permutations that have an identity block of size at least cn is given by

$$\nu_c \leq \binom{n}{cn}((1-c)n)! = \frac{n!}{(cn)!}.$$

Thus,

$$\mathbb{P}\left[|\mathcal{I}_\pi| \geq cn\right] \leq \frac{1}{\sqrt{2\pi}} \exp\left(-(cn + \frac{1}{2})\log(cn) + cn\right),$$

from Stirling's approximation.

Lengths and number of cycles in a random permutation may be analyzed as detailed in [20]. In particular, we note that for a random permutation π,

$$\mathbb{E}\left[\kappa_\pi\right] = \log n + O(1).$$

Using Markov's Inequality, the result follows. ■

Following the observations of Lemma 1, we assume that for any $\pi \in \Pi$, $\kappa_\pi = o\left(\frac{n}{\log(n)}\right)$, i.e., the number of permutation cycles does not grow very fast. Secondly, let $|\mathcal{I}_\pi| \leq \gamma_\pi n$ for any $\pi \in \Pi$, for $\gamma_\pi = o(1)$.

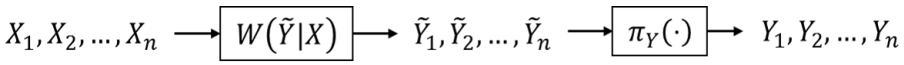

Fig. 2. Model of the two image registration problem: Image $\mathbf{Y}$ is to be registered with source image $\mathbf{X}$.

C. Performance Metrics

We now formally introduce performance metrics to quantify performance of the joint clustering and registration algorithms.

Definition 5 (Correct Clustering): A *clustering* of a set of images $\{\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(m)}\}$ is a partition P of $[m]$. The sets of a partition are referred to as *clusters*. The clustering is said to be *correct* if

$$i, j \in C \Leftrightarrow \mathbf{R}^{(i)} = \mathbf{R}^{(j)}, \text{ for all } i, j \in [m], \ C \in P.$$

Let $\mathcal{P}$ be the set of all partitions of $[m]$. For a given collection of scenes, we shall represent the correct clustering by P^*.

Definition 6 (Partition Ordering): A partition P is *finer* than P', if the following ordering holds:

$$P \preceq P' \Leftrightarrow \text{ for all } C \in P, \text{ there exists } C' \in P' : C \subseteq C'.$$

Similarly, a partition P is said to be *denser* than P', $P \succeq P'$, if $P' \preceq P$.

Definition 7 (Correct Registration): The *correct registration* of an image $\mathbf{X}$ transformed by $\pi \in \Pi$ is $\hat{\pi} = \pi^{-1}$.

Definition 8 (Universal Clustering and Registration Algorithm): A *universal clustering and registration algorithm* is a sequence of functions $\Phi^{(n)} : \left\{\mathbf{X}^{(1)}, \ldots, \mathbf{X}^{(m)}\right\} \to \mathcal{P} \times \Pi^m$ that are designed in the absence of knowledge of W, $\{p_1, \ldots, p_\ell\}$, and P_R. Here the index n corresponds to the number of pixels in each image.

In this work, we focus on the $0-1$ loss function to quantify the performance of algorithms.

Definition 9 (Error Probability): The *error probability* of an algorithm $\Phi^{(n)}$ that outputs $\hat{P} \in \mathcal{P}$ and $(\hat{\pi}_1, \ldots, \hat{\pi}_m) \in \Pi^m$ is

$$P_e(\Phi^{(n)}) = \mathbb{P}\left[\{\hat{P} \neq P^*\} \cup \{\text{there exists } i \in [m] : \hat{\pi}_i \neq \pi_i^{-1}\}\right]. \tag{3}$$

Definition 10 (Asymptotic Consistency): A sequence of decoders $\Phi^{(n)}$ is said to be *asymptotically consistent* if

$$\lim_{n\to\infty} P_e(\Phi^{(n)}) = 0.$$

In particular, it is *exponentially consistent* if

$$\lim_{n\to\infty} -\log P_e(\Phi^{(n)}) > 0.$$

Definition 11 (Error Exponent): The *error exponent* of an image clustering and registration algorithm $\Phi^{(n)}(\cdot)$ is

$$\mathcal{E}(\Phi^{(n)}) = \lim_{n\to\infty} -\frac{1}{n} \log P_e(\Phi^{(n)}). \tag{4}$$

For simplicity, we will use Φ to denote $\Phi^{(n)}$ when it is clear from context.

III. Registration of two images

We now restrict the problem to that of registering two images. Specifically, we consider the scenario with $m = 2$, $\ell = 1$. Thus the problem reduces to one of registering an image $\mathbf{Y}$ obtained as a result of transforming the output of an equivalent discrete memoryless channel W, given input image (reference) $\mathbf{X}$. The model of the channel is depicted in Fig. 2.

This problem has been well-studied in practice and a popular heuristic method used extensively is the maximum mutual information (MMI) method defined as

$$\hat{\pi}_{\text{MMI}} = \arg\max_{\pi\in\Pi} \hat{I}(X; Y_\pi), \tag{5}$$

where $\hat{I}(X;Y)$ is the mutual information corresponding to the empirical joint distribution

$$\hat{P}(x,y) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\left\{X_i = x, Y_i = y\right\}.,$$

where $\mathbf{1}\left\{\cdot\right\}$ is the indicator function. It is worth noting that the MMI method for image registration is universal.

Since transformations are chosen uniformly at random, the maximum likelihood (ML) estimate is Bayes optimal:

$$\hat{\pi}_{\text{ML}} = \arg\max_{\pi\in\Pi} \prod_{i=1}^{n} W\left(Y_{\pi(i)}|X_i\right). \tag{6}$$

We first show that the MMI method, and consequently the ML method, are exponentially consistent. We then show that the error exponent of MMI matches that of ML.

A. Exponential Consistency

The empirical mutual information of i.i.d. samples is asymptotically exponentially consistent.

Theorem 1: Let $\{(X_1,Y_1),\ldots,(X_n,Y_n)\}$ be random samples drawn i.i.d. according to the joint distribution p defined on the finite space $\mathcal{X}\times\mathcal{Y}$. For fixed alphabet sizes, $|\mathcal{X}|, |\mathcal{Y}|$, the ML estimate of entropy and mutual information are asymptotically exponentially consistent and satisfy

$$\mathbb{P}\left[|\hat{H}(X) - H(X)| > \epsilon\right] \leq (n+1)^{|\mathcal{X}|}\exp\left(-cn\epsilon^4\right), \tag{7}$$

$$\mathbb{P}\left[|\hat{I}(X;Y) - I(X;Y)| > \epsilon\right] \leq 3(n+1)^{|\mathcal{X}||\mathcal{Y}|}\exp\left(-\tilde{c}n\epsilon^4\right), \tag{8}$$

where $c = \left(2|\mathcal{X}|^2\log 2\right)^{-1}$ and $\tilde{c} = \left(32\max\{|\mathcal{X}|,|\mathcal{Y}|\}^2\log 2\right)^{-1}$.
The proof is given in [5, Lemma 7].

Theorem 2: The MMI and ML image registration estimates are exponentially consistent.

Proof: Let $\Phi_{\text{MMI}}(\mathbf{X},\mathbf{Y}) = \hat{\pi}_{\text{MMI}}$ and let the correct registration be π^*. Then,

$$\begin{aligned}
P_e(\Phi_{\text{MMI}}) &= \mathbb{P}\left[\hat{\pi}_{\text{MMI}} \neq \pi^*\right] \\
&\leq \sum_{\pi\in\Pi} \mathbb{P}\left[\hat{I}(X;Y_\pi) > \hat{I}(X;\tilde{Y})\right] &(9)\\
&\leq \sum_{\pi\in\Pi} \mathbb{P}\left[|\hat{I}(X;Y_\pi) - (\hat{I}(X;\tilde{Y}) - I(X;\tilde{Y}))| > I(X;\tilde{Y})\right] \\
&\leq \sum_{\pi\in\Pi} \mathbb{P}\left[\hat{I}(X;Y_\pi) + |\hat{I}(X;\tilde{Y}) - I(X;\tilde{Y})| > I(X;\tilde{Y})\right] &(10)\\
&\leq 2|\Pi|\exp\left\{-\tilde{c}n(I(X;\tilde{Y}))^4\right\}, &(11)
\end{aligned}$$

where (9), (10), and (11) follow from the union bound, the triangle inequality, and (8), respectively.

Thus MMI is exponentially consistent as $|\Pi| = O(n^\alpha)$. Finally, $P_e(\Phi_{\text{ML}}) \leq P_e(\Phi_{\text{MMI}})$ and thus, the ML estimate is also exponentially consistent. ∎

Theorem 2 also proves that there exists $\epsilon > 0$ such that $\mathcal{E}(\Phi_{\text{ML}}) \geq \mathcal{E}(\Phi_{\text{MMI}}) \geq \epsilon$.

B. Whittle's Law and Markov Types

We now summarize a few results on the number of types and Markov types which are eventually used to analyze the error exponent of image registration.

Consider a sequence $\mathbf{x}\in\mathcal{X}^n$. The empirical distribution q_X of $\mathbf{x}$ is referred to as the *type* of the sequence. Let $X\sim q_X$ be a dummy random variable. Let T_X^n be the set of all sequences of length n, of type q_X. The number of possible types of sequences of length n is polynomial in n, i.e., $O(n^{|\mathcal{X}|})$ [21].

The number of sequences of length n, of type q_X, is

$$|T_X^n| = \frac{n!}{\prod_{a\in\mathcal{X}}(nq_X(a))!}.$$

From the bounds on multinomial coefficients, the number of sequences of length n and type q is bounded as [21]

$$(n+1)^{-|\mathcal{X}|}2^{nH(X)} \leq |T_X^n| \leq 2^{nH(X)}. \tag{12}$$

Consider a Markov chain defined on the space $[k]$. Given a sequence of $n+1$ samples from $[k]$ we can compute the matrix $\mathbf{F}$ of transition counts, where F_{ij} corresponds to the number of transitions from state i to state j. By Whittle's formula [22], the number of sequences $(a_1, \ldots, a_{n+1})$ with $a_i \in [k], i \in [n+1]$, with $a_1 = u$ and $a_{n+1} = v$ is given by

$$N_{uv}^{(n)}(F) = \prod_{i\in[k]} \frac{(\sum_{j\in[k]} F_{ij})!}{\prod_{j\in[k]} F_{ij}!} G_{vu}^*, \tag{13}$$

where G_{vu}^* corresponds to the (v,u)th cofactor of the matrix $G = \{g_{ij}\}_{i,j\in[k]}$ with

$$g_{ij} = \mathbf{1}\{i=j\} - \frac{F_{ij}}{\sum_{j\in[k]} F_{ij}}.$$

The first-order Markov type of a sequence $\mathbf{x} \in \mathcal{X}^n$ is defined as the empirical distribution q_{X_0,X_1}, given by

$$q_{X_0,X_1}(a_0,a_1) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{(x_i, x_{i+1}) = (a_0,a_1)\}.$$

Here we assume that the sequence is cyclic with period n, i.e., for any $i > 0, x_{n+i} = x_i$. Let $(X_0, X_1) \sim q_{X_0,X_1}$. Then, from (13), the set of sequences of type q_{X_0,X_1}, T_{X_0,X_1}^n, satisfies

$$|T_{X_0,X_1}^n| = \left(\sum_{a\in\mathcal{X}} G_{a,a}^*\right) \prod_{a_0\in\mathcal{X}} \frac{(nq_0(a_0))!}{\prod_{a_1\in\mathcal{X}}(nq_{\mathbf{x}_0,\mathbf{x}_1}(a_0,a_1))!}.$$

From the definition of G, we can bound the trace of the cofactor matrix of G as

$$\frac{|\mathcal{X}|}{(n+1)^{|\mathcal{X}|}} \leq \sum_{a\in\mathcal{X}} G_{a,a}^* \leq |\mathcal{X}|.$$

Again using the bounds on multinomial coefficients, we have

$$|\mathcal{X}|(n+1)^{-(|\mathcal{X}|^2+|\mathcal{X}|)}2^{n(H(X_0,X_1)-H(X_0))} \leq |T_{X_0,X_1}^n| \leq |\mathcal{X}|2^{n(H(X_0,X_1)-H(X_0))}. \tag{14}$$

The joint first-order Markov type of a pair of sequences $\mathbf{x} \in \mathcal{X}^n$, $\mathbf{y} \in \mathcal{Y}^n$ is the empirical distribution

$$q_{X_0,X_1,Y}(a_0,a_1,b) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{(x_i, x_{i+1}, y_i) = (a_0,a_1,b)\}.$$

Then given $\mathbf{x}$, the set of conditional first-order Markov type sequences, $T_{Y|X_0,X_1}^n(\mathbf{x})$ satisfies [18]

$$(n+1)^{-|\mathcal{X}|^2|\mathcal{Y}|}2^{n(H(X_0,X_1,Y)-H(X_0,X_1))} \leq |T_{Y|X_0,X_1}^n(\mathbf{x})| \leq 2^{n(H(X_0,X_1,Y)-H(X_0,X_1))}. \tag{15}$$

Lemma 2: Let $\pi_1, \pi_2 \in \Pi$ be any two non-overlapping permutations and let $\mathbf{x}_{\pi_1}, \mathbf{x}_{\pi_2}$ be the corresponding permutations of $\mathbf{x}$. Let $\pi_1^{-1} \circ \pi_2 \in \Pi$ be a simple permutation. Then, for every $\mathbf{x}$, there exists $\tilde{\mathbf{x}}$, such that

$$|T_{X_{\pi_1},X_{\pi_2}}^n| = |T_{X_0,X_1}^n|, \quad |T_{Y|X_{\pi_1},X_{\pi_2}}^n(\mathbf{x})| = |T_{Y|X_0,X_1}^n(\tilde{\mathbf{x}})|.$$

Proof: Since the permutations are non-overlapping, there exists a bijection from $T_{X_{\pi_1},X_{\pi_2}}^n$ to T_{X_0,X_1}^n, where $(X_0,X_1) \sim q_{X_{\pi_1},X_{\pi_2}}$. Specifically, consider the permutation $\pi \in \Pi$ defined iteratively as $\pi(i+1) = \pi_2(\pi_1^{-1}(\pi(i)))$, with $\pi(1) = \pi_1(1)$. Then, for any $\mathbf{x} \in T_{X_{\pi_1},X_{\pi_2}}^n$, the sequence $\mathbf{x}_\pi \in T_{X_0,X_1}^n$. Further, this map is invertible and thus the sets are of equal size.

The result for the conditional types also follows from the same argument. ∎

In particular, Lemma 2 implies that $|T_{X_{\pi_1},X_{\pi_2}}^n|$ and $|T_{Y|X_{\pi_1},X_{\pi_2}}^n(\mathbf{x})|$, satisfy (14) and (15) respectively. We now show that the result of Lemma 2 can be extended to any two permutations $\pi_1, \pi_2 \in \Pi$.

Lemma 3: Let $\pi_1, \pi_2 \in \Pi$. For any $\mathbf{x}$,

$$|T_{X_{\pi_1},X_{\pi_2}}^n| = 2^{n\left(H(q_{X_{\pi_1},X_{\pi_2}})-H(q_X)+o(1)\right)}.$$

Proof: Let $\pi = \pi_1^{-1} \circ \pi_2$ and $\kappa = \kappa_\pi$. For $i \in [\kappa]$ let the length of permutation cycle i of π be $\alpha_i n$ for $\alpha_i \in (0, 1]$. Further, $\sum_{i=1}^{\kappa} \alpha_i \leq 1$. Let $\mathcal{I}_\pi$ be the identity block of π and let $\gamma = \gamma_\pi$. Then we have the decomposition

$$q_{X_{\pi_1}, X_{\pi_2}}(a_0, a_1) = \sum_{i=1}^{\kappa} \alpha_i q_i(a_0, a_1) + \gamma q_{\mathcal{I}}(a_0, a_1), \tag{16}$$

for all $(a_0, a_1) \in \mathcal{X}^2$. Here, q_i is the first-order Markov type defined on the ith permutation cycle of π and $q_{\mathcal{I}}$ is the zeroth-order Markov type corresponding to the identity block of π.

From (16), we see that given a valid decomposition of types $\{q_i\}, q_{\mathcal{I}}$, the number of sequences can be computed as a product of the number of subsequences of each type. That is,

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| = \sum |T^{\gamma n}_{q_{\mathcal{I}}}| \prod_{i=1}^{\kappa} |T^{\alpha_i n}_{q_i}|,$$

where the summation is over all valid decompositions given in (16).

Additionally, from Lemma 2 we know the number of valid subsequences of each type. Let q_i' be the marginal corresponding to the first-order Markov type q_i.

Thus, we upper bound the size of the set as

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| \leq \sum 2^{\gamma n H(q_{\mathcal{I}})} \prod_{i=1}^{\kappa} |\mathcal{X}| 2^{\alpha_i n (H(q_i) - H(q_i'))} \tag{17}$$

$$\leq |\mathcal{X}|^{\kappa} (\gamma n + 1)^{|\mathcal{X}|} \prod_{i=1}^{\kappa} (\alpha_i n + 1)^{|\mathcal{X}|^2} 2^{n M(q_{X_{\pi_1}, X_{\pi_2}})}, \tag{18}$$

where

$$M(q_{X_{\pi_1}, X_{\pi_2}}) = \max \gamma H(q_{\mathcal{I}}) + \sum_{i=1}^{\kappa} \alpha_i (H(q_i) - H(q_i')),$$

the maximum taken over all valid decompositions defined by (16). Here, (17) follows from (12) and (14), and (17) follows from the fact that the total number of possible types is polynomial in the length of the sequence. Since $\kappa = o\left(\frac{n}{\log(n)}\right)$,

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| \leq 2^{n\left(M(q_{X_{\pi_1}, X_{\pi_2}}) + o(1)\right)}.$$

Now, let $q_i''(a_0, a_1) = \frac{1}{|\mathcal{X}|} q_i'(a_1)$, for all $i \in [\kappa], (a_0, a_1) \in \mathcal{X}^2$. And let $q = q_{X_{\pi_1}, X_{\pi_2}}$. Since $\gamma = o(1)$, using Jensen's inequality, we have

$$\begin{aligned} \sum_{i=1}^{\kappa} \alpha_i (H(q_i) - H(q_i')) &= \sum_{i=1}^{\kappa} \alpha_i \left[\log(|\mathcal{X}|) - D\left(q_i \| q_i''\right)\right] \\ &\leq \log(|\mathcal{X}|) - D\left(q \| q''\right) \\ &= H(q) - H(q'). \end{aligned}$$

Thus,

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| \leq 2^{n(H(q) - H(q') + o(1))}.$$

To obtain the lower bound, we note that the total number of sequences is at least the number of sequences obtained from any single valid decomposition. Thus from (12) and (14),

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| \geq 2^{n\left(M(q_{X_{\pi_1}, X_{\pi_2}}) + o(1)\right)}.$$

Now, for large n, consider $S = \{i \in \kappa : \alpha_i n = \Omega(n^\beta)\}$ for some $\beta > 0$. Any other cycle of smaller length contributes $o(1)$ to the exponent owing to Lemma 1. One viable decomposition of (16) is to have $q_i = q$, for all $i \in [\kappa]$. However, the lengths of the subsequences are different and q may not be a valid type of the corresponding length. Nevertheless, for each $i \in S$, there exists a type q_i such that

$$d_{\mathrm{TV}}(q_i, q) \leq \frac{|\mathcal{X}|}{2\alpha_i n},$$

where $d_{\mathrm{TV}}(\cdot)$ is the total variational distance. Further, entropy is continuous in the distribution [23] and satisfies

$$|H(q_i) - H(q)| \leq \frac{|\mathcal{X}|}{\alpha_i n} \log(\alpha_i n) = o(1).$$

This in turn indicates that

$$|T^n_{X_{\pi_1}, X_{\pi_2}}| \geq 2^{n(H(q)-H(q')+o(1))}.$$

Hence the result follows. ∎

A similar decomposition rule follows for the conditional type as well.

C. Error Analysis

We are interested in the error exponent of the MMI-based image registration, in comparison to ML. We first note that the error exponent of the problem is characterized by the pair of transformations that are the hardest to compare.

Define $\Psi_{\pi,\pi'}$ as the binary hypothesis testing problem corresponding to image registration when the allowed transformations are only $\{\pi, \pi'\}$. Let $P_{\pi,\pi'}(\Phi), \mathcal{E}_{\pi,\pi'}(\Phi)$ be the corresponding error probability and error exponent.

Lemma 4: Let Φ be an asymptotically exponentially consistent estimator. Then,

$$\mathcal{E}(\Phi) = \min_{\pi,\pi' \in \Pi} \mathcal{E}_{\pi,\pi'}(\Phi). \tag{19}$$

Proof: Let $\hat{\pi}$ be the estimate output by Φ and $\pi^* \sim \mathrm{Unif}(\Pi)$ be the correct registration. We first upper bound the error probability as

$$\begin{aligned} P_e(\Phi) &= \mathbb{P}\left[\hat{\pi} \neq \pi^*\right] \\ &\leq \sum_{\pi,\pi' \in \Pi, \pi \neq \pi'} \frac{1}{|\Pi|} \mathbb{P}\left[\hat{\pi} = \pi | \pi^* = \pi'\right] \qquad (20) \\ &\leq \frac{2}{|\Pi|} \sum_{\pi,\pi' \in \Pi, \pi \neq \pi'} P_{\pi,\pi'}(\Phi), \end{aligned}$$

where (20) follows from the union bound.

Additionally, we have

$$\begin{aligned} P_e(\Phi) &= \frac{1}{|\Pi|} \sum_{\pi \in \Pi} \sum_{\pi' \in \Pi \setminus \{\pi\}} \mathbb{P}\left[\hat{\pi} = \pi' | \pi^* = \pi\right] \\ &\geq \frac{1}{|\Pi|} \max_{\pi,\pi' \in \Pi} P_{\pi,\pi'}(\Phi). \end{aligned}$$

Finally, since $|\Pi| = O(n^\alpha)$, the result follows. ∎

Thus, it suffices to consider the binary hypothesis tests to study the error exponent of image registration.

Theorem 3: Let $\{\pi_1, \pi_2\} \subseteq \Pi$. Then,

$$\lim_{n \to \infty} \frac{P_{\pi_1,\pi_2}(\Phi_{\mathrm{MMI}})}{P_{\pi_1,\pi_2}(\Phi_{\mathrm{ML}})} = 1. \tag{21}$$

Proof: Probability of i.i.d. sequences are defined by their joint type. Thus, we have

$$\begin{aligned} P_{\pi_1,\pi_2}(\Phi_{\mathrm{MMI}}) &= \mathbb{P}\left[\hat{\pi}_{\mathrm{MMI}} \neq \pi^*\right] \\ &= \sum_{\mathbf{x} \in \mathcal{X}^n} \sum_{\mathbf{y} \in \mathcal{Y}^n} \mathbb{P}\left[\mathbf{x}, \mathbf{y}\right] \mathbf{1}\left\{\hat{\pi}_{\mathrm{MMI}} \neq \pi^*\right\} \\ &= \sum_{q} \left(\prod_{a \in \mathcal{X}} \prod_{b \in \mathcal{Y}} \left(\mathbb{P}\left[a, b\right]\right)^{nq(a,b)} \right) \nu_{\mathrm{MMI}}(q), \qquad (22) \end{aligned}$$

where the summation in (22) is over the set of all joint types of sequences of length n and $\nu_{\mathrm{MMI}}(q)$ is the number of sequences $(\mathbf{x}, \mathbf{y})$ of length n with joint type q such that the MMI algorithm makes a decision error.

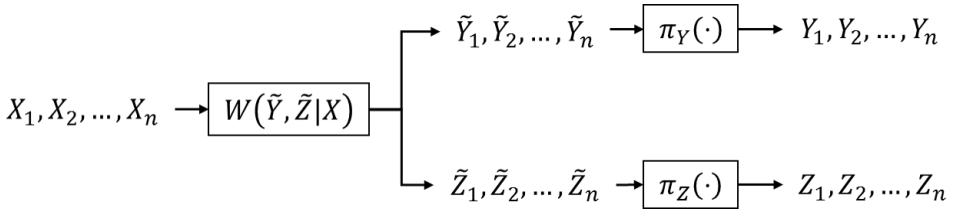

Fig. 3. Model of the three-image registration problem: Images $\mathbf{Y}, \mathbf{Z}$ are to be registered with source image $\mathbf{X}$.

If a sequence $\mathbf{y} \in T^n_{Y|X_{\pi_1},X_{\pi_2}}(\mathbf{x})$ is in error, then all sequences in $T^n_{Y|X_{\pi_1},X_{\pi_2}}(\mathbf{x})$ are in error as $\mathbf{x}$ is i.i.d.. Now, we can decompose the number of sequences under error as follows

$$\begin{aligned}\nu(q) &= \sum_{\mathbf{x}\in\mathcal{X}^n} \sum_{T^n_{Y|X_{\pi_1},X_{\pi_2}} \subseteq T^n_{Y|X}:\text{error}} |T^n_{Y|X_{\pi_1},X_{\pi_2}}(\mathbf{x})| \\ &= \sum_{T^n_{X_{\pi_1},X_{\pi_2}} \subseteq T^n_X} |T_{X_{\pi_1},X_{\pi_2}}| \sum |T^n_{Y|X_{\pi_1},X_{\pi_2}}|,\end{aligned}$$

where the sum is taken over $T^n_{Y|X_{\pi_1},X_{\pi_2}} \subseteq T^n_{Y|X}$ such that there is a decision error. The final line follows from the fact that given the joint first-order Markov type, the size of the conditional type is independent of the exact sequence $\mathbf{x}$.

Finally, we have

$$\begin{aligned}P_{\pi_1,\pi_2}(\Phi_{\text{MMI}}) &= \sum_q \prod_{a\in\mathcal{X}} \prod_{b\in\mathcal{Y}} (\mathbb{P}[a,b])^{nq(a,b)} \nu_{\text{ML}}(q) \left\{ \frac{\nu_{\text{MMI}}(q)}{\nu_{\text{ML}}(q)} \right\} \\ &\leq P_{\pi_1,\pi_2}(\Phi_{\text{ML}}) \max_q \left\{ \frac{\nu_{\text{MMI}}(q)}{\nu_{\text{ML}}(q)} \right\}. \end{aligned} \tag{23}$$

The result follows from Lemma 5. ■

Lemma 5:

$$\lim_{n\to\infty} \max_q \left\{ \frac{\nu_{\text{MMI}}(q)}{\nu_{\text{ML}}(q)} \right\} = 1. \tag{24}$$

Proof: We first observe that for images with i.i.d. pixels, image registration using MMI is the same as minimizing the joint entropy. That is,

$$\begin{aligned}\max_{\pi\in\Pi} \hat{I}(X;\pi(Y)) &= \max_{\pi\in\Pi} \hat{H}(X) + \hat{H}(\pi(Y)) - \hat{H}(X,\pi(Y)) \\ &= \hat{H}(X) + \hat{H}(\tilde{Y}) - \min_{\pi\in\Pi} \hat{H}(X,\pi(Y)).\end{aligned}$$

Additionally we know that there exists a bijective mapping from sequences corresponding to the permutations and the sequences of the corresponding first-order Markov type from Lemmas 2 and 3. Thus, the result follows from [18, Lemma 1]. ■

Theorem 4:

$$\mathcal{E}(\Phi_{\text{MMI}}) = \mathcal{E}(\Phi_{\text{ML}}). \tag{25}$$

Proof: This follows as a direct consequence of Theorem 3 and Lemma 4. ■

Thus, we can see that the using MMI for image registration is not only universal, but also asymptotically optimal. We next study the problem of image registration for multiple images.

IV. Multi-image Registration

Having studied the problem of registering two images in a universal sense, we now consider the problem of aligning multiple copies of the same image. For simplicity, let us consider the task of aligning three images. The results obtained for this case can directly be extended to any finite number of copies of images to be aligned. Again, let $\mathbf{X}$ be the source image and $\mathbf{Y}, \mathbf{Z}$ be the noisy, transformed versions to be aligned as shown in Fig. 3.

Here, the ML estimates are

$$(\hat{\pi}_{Y,ML}, \hat{\pi}_{Z,ML}) = \arg \max_{(\pi_1,\pi_2)\in\Pi^2} \prod_{i=1}^{n} W\left(Y_{\pi_1(i)}, Z_{\pi_2(i)}|X_i\right). \tag{26}$$

A. MMI is not optimal

We know that MMI is asymptotically optimal for aligning two images. Thus the natural question is if performing MMI in a pairwise sense is optimal for multi-image registration. That is,

$$\hat{\pi}_Y = \arg\max_{\pi\in\Pi} \hat{I}(X;\pi(Y)), \quad \hat{\pi}_Z = \arg\max_{\pi\in\Pi} \hat{I}(X;\pi(Z)). \tag{27}$$

We first address this question, making it explicit that pairwise MMI method is sub-optimal even though individual transformations are chosen independently and uniformly from Π.

Theorem 5: There exists a channel W and a source distribution, such that pairwise MMI method for multi-image registration is sub-optimal.

Proof: Let $X_i \overset{\text{i.i.d.}}{\sim} \text{Bern}(1/2), i \in [n]$. Consider the physically degraded images $\mathbf{Y}, \mathbf{Z}$ obtained as the output of the channel defined by

$$W(y,z|x) = W_1(y|x)\, W_2(z|y).$$

This scenario is depicted in Fig. 4.

Naturally, the ML estimate of the transformations is obtained by registering image $\mathbf{Z}$ to $\mathbf{Y}$ and subsequently registering $\mathbf{Y}$ to $\mathbf{X}$, instead of registering each of the images in a pairwise sense to $\mathbf{X}$. That is, from (26)

$$(\hat{\pi}_{Y,\text{ML}}, \hat{\pi}_{Z,\text{ML}}) = \arg \max_{(\pi_1,\pi_2)\in\Pi^2} \prod_{i=1}^{n} W_1(Y_{\pi_1(i)}|X_i) W_2(Z_{\pi_2(i)}|Y_{\pi_1(i)}).$$

It is evident that π_Z is estimated based on the estimate of π_Y.

Let $\mathcal{E}_W(\Phi_{\text{ML}})$ be the error exponent of Φ_{ML} for the physically degraded channel. Since the ML estimate here involves registration of $\mathbf{Y}$ to $\mathbf{X}$ and $\mathbf{Z}$ to $\mathbf{Y}$, the effective error exponent is given by

$$\mathcal{E}_W(\Phi_{\text{ML}}) = \min\left\{\mathcal{E}_{W_1}(\Phi_{\text{ML}}), \mathcal{E}_{W_2}(\Phi_{\text{ML}})\right\}.$$

Let $\mathcal{E}_Q(\Phi_{\text{MMI}})$ be the error exponent of image registration using MMI for the channel Q. Then, the error exponent of pairwise MMI based image registration is given by

$$\mathcal{E}(\Phi_{\text{MMI}}) = \min\left\{\mathcal{E}_{W_1}(\Phi_{\text{MMI}}), \mathcal{E}_{W_1 * W_2}(\Phi_{\text{MMI}})\right\}.$$

We know that for registration of two images, MMI is asymptotically optimal and thus

$$\mathcal{E}_{W_1}(\Phi_{\text{MMI}}) = \mathcal{E}_{W_1}(\Phi_{\text{ML}}), \quad \mathcal{E}_{W_1 * W_2}(\Phi_{\text{MMI}}) = \mathcal{E}_{W_1 * W_2}(\Phi_{\text{ML}}).$$

More specifically, let $W_1 = BSC(\alpha)$ and $W_2 = BSC(\beta)$ for some $\alpha, \beta \in (0, 1/2)$. Let, $W_1 * W_2 = BSC(\gamma)$, where $\gamma = \alpha(1-\beta) + (1-\alpha)\beta > \max\{\alpha, \beta\}$. Then,

$$\begin{aligned}
\mathcal{E}(\Phi_{\text{MMI}}) &\leq \mathcal{E}_{W_1 * W_2}(\Phi_{\text{MMI}}) \\
&= \mathcal{E}_{W_1 * W_2}(\Phi_{\text{ML}}) \\
&< \min\left\{\mathcal{E}_{W_1}(\Phi_{\text{ML}}), \mathcal{E}_{W_2}(\Phi_{\text{ML}})\right\} \\
&= \mathcal{E}_W(\Phi_{\text{ML}}).
\end{aligned} \tag{28}$$

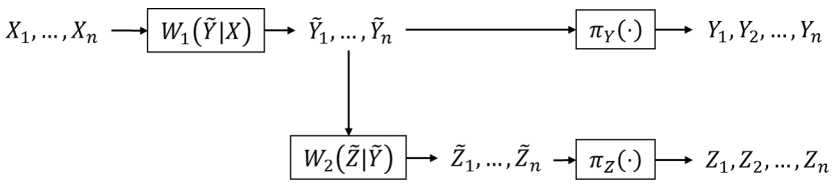

Fig. 4. three-image registration problem with image $\mathbf{Z}$ a physically degraded version of $\mathbf{Y}$ and to be registered with source image $\mathbf{X}$.

Hence the result follows. ■

We note here that the sub-optimality of the pairwise MMI method arises from the rigidity of the scheme, i.e., specifically in the underlying principle that does not take into account the fact that images $\mathbf{Y}$ and $\mathbf{Z}$ are dependent conditioned on $\mathbf{X}$. Thus, it would be prudent to design a universal registration algorithm that takes into account such dependencies.

B. Max Multiinformation based multi-image registration

To take into consideration all correlations across images, we make use of the multiinformation functional [24] which is a more inclusive formulation of the underlying information in the system.

Definition 12 (Multiinformation): The *multiinformation* of random variables $X_1, \ldots, X_n$ is defined as

$$I_M(X_1; \ldots; X_n) = \sum_{i=1}^{n} H(X_i) - H(X_1, \ldots, X_n). \tag{29}$$

The chain rule for multiinformation is given by

$$I_M(X_1; \ldots; X_n) = \sum_{i=2}^{n} I(X_i; X^{(i-1)}), \tag{30}$$

where $X^{(i)} = (X_1, \ldots, X_i)$.

Let us define a new multi-image registration method, which we refer to as the maximum multiinformation (MM) estimate, defined for the three-image case as

$$(\hat{\pi}_{Y,MM}, \hat{\pi}_{Z,MM}) = \arg \max_{(\pi_1, \pi_2) \in \Pi^2} \hat{I}_M(X; Y_{\pi_1}; Z_{\pi_2}), \tag{31}$$

where $\hat{I}_M(\cdot)$ is the empirical estimate of multiinformation. The estimator can correspondingly be generalized to any finite number of images. It is worth noting that the estimates obtained are the same as that of the pairwise MMI method, when $\mathbf{Y}$ and $\mathbf{Z}$ are conditionally independent (empirically) given $\mathbf{X}$.

Lemma 6: The MM estimates are exponentially consistent for multi-image registration.

Proof: The proof is analogous to that of Theorem 2, and follows from the union bound and Theorem 1. ■

We now show that the MM method is asymptotically optimal in the error exponent. Again we make use of a type counting argument.

C. Error Analysis

Again, we compare the error exponent of MM with respect to that of ML. We again show that the error exponent is characterized by the pair of transformations of $\mathbf{Y}, \mathbf{Z}$ that are the hardest to differentiate.

Let $\bar{\pi}, \bar{\pi}' \in \Pi^2$ and let $\psi(\bar{\pi}, \bar{\pi}')$ be the binary hypothesis test of the problem of three-image registration when the set of permitted transformations is restricted to $\{\bar{\pi}, \bar{\pi}'\}$. Let $P_{\bar{\pi},\bar{\pi}'}(\Phi), \mathcal{E}_{\bar{\pi},\bar{\pi}'}(\Phi)$ be the error probability and the error exponent of Φ respectively.

Lemma 7: Let Φ be an asymptotically exponentially consistent estimator of three-image registration. Then,

$$\mathcal{E}(\Phi) = \min_{\bar{\pi},\bar{\pi}' \in \Pi^2} \mathcal{E}_{\bar{\pi},\bar{\pi}'}(\Phi). \tag{32}$$

Proof: The lower bound on the error exponent follows directly from the union bound. Thus,

$$\mathcal{E}(\Phi) \geq \min_{\bar{\pi},\bar{\pi}' \in \Pi^2} \mathcal{E}_{\bar{\pi},\bar{\pi}'}(\Phi).$$

On the other hand, we have

$$\begin{aligned} P_e(\Phi) &= \frac{1}{|\Pi|^2} \sum_{\bar{\pi} \in \Pi} \sum_{\bar{\pi}' \in \Pi^2 \setminus \bar{\pi}} \mathbb{P}\left[\hat{\pi} = \bar{\pi}' | \bar{\pi}^* = \bar{\pi}\right] \\ &\geq \frac{1}{|\Pi|^2} \max_{\bar{\pi},\bar{\pi}' \in \Pi^2} P_{\bar{\pi},\bar{\pi}'}(\Phi). \end{aligned}$$

Thus, the result follows. ■

Let $\bar{\pi}_0 = (\pi_0, \pi_0)$ correspond to the scenario where the image is not transformed.

Lemma 8: For any $\bar{\pi}_1, \bar{\pi}_2 \in \Pi^2$, with $\bar{\pi}_i = (\pi_i, \pi_i')$, there exists $\bar{\pi} = (\pi, \pi)$ such that

$$\mathcal{E}_{\bar{\pi}_1,\bar{\pi}_2}(\Phi) \geq \mathcal{E}_{\bar{\pi}_0,\bar{\pi}}(\Phi),$$

for $\Phi \in \{\Phi_{\text{ML}}, \Phi_{\text{MM}}\}$.

Proof: We prove the result for the ML decoder. The proof can be extended directly to the MM decoder as well.

Let $\tilde{\pi}_1$ be the permutation such that $\tilde{\pi}_1 \circ \pi_1' = \pi_1$. That is, the application of the transformation $\tilde{\pi}_1$ to an image that has been transformed by π_1' results in an image that is effectively transformed by π_1. Let $\tilde{\pi}_2 = \tilde{\pi}_1 \circ \pi_2'$.

To obtain the ML decision for $\psi(\bar{\pi}_1, \bar{\pi}_2)$, let

$$(\hat{\pi}_Y, \hat{\pi}_Z) = \arg \max_{(\pi_Y,\pi_Z)\in\{(\pi_1,\pi_1),(\pi_2,\tilde{\pi}_2)\}} W(\mathbf{Y}_{\pi_Y}, \mathbf{Z}'_{\pi_Z}|\mathbf{X}),$$

where $\mathbf{Z}' = \mathbf{Z}_{\tilde{\pi}_1}$, is the received image, transformed by $\tilde{\pi}_1$. Then, $(\hat{\pi}_{Y,\text{ML}}, \hat{\pi}_{Z,\text{ML}}) = (\hat{\pi}_Y, \tilde{\pi}_1^{-1} \circ \hat{\pi}_Z)$.

Since the source is i.i.d. and the channel memoryless, this in turn indicates that

$$\mathcal{E}_{\bar{\pi}_1,\bar{\pi}_2}(\Phi_{\text{ML}}) = \mathcal{E}_{(\pi_1,\pi_1),(\pi_2,\tilde{\pi}_2)}(\Phi_{\text{ML}}).$$

Finally, if π_1^{-1} is the inverse of the transformation π_1, let $\pi = \pi_1^{-1} \circ \pi_2, \pi' = \pi_1^{-1} \circ \tilde{\pi}_2$. Then,

$$(\hat{\pi}_Y, \hat{\pi}_Z) = \arg \max_{(\pi_Y,\pi_Z)\in\{(\pi_0,\pi_0),(\pi,\pi')\}} W(\mathbf{Y}'_{\pi_Y}, \mathbf{Z}'_{\pi_Z}|\mathbf{X}),$$

where $\mathbf{Y}' = \mathbf{Y}_{\pi_1^{-1}}$ and $\mathbf{Z}' = \mathbf{Z}_{\pi_1^{-1}\circ\tilde{\pi}_1}$, then $(\hat{\pi}_{Y,\text{ML}}, \hat{\pi}_{Z,\text{ML}}) = (\pi_1 \circ \hat{\pi}_Y, (\pi_1^{-1} \circ \tilde{\pi}_1)^{-1} \circ \hat{\pi}_Z)$.

We note now that

$$D(P_{(\pi_0,\pi_0)}(X,Y,Z)\|P_{(\pi,\pi')}(X,Y,Z)) = I_M(X;Y;Z). \tag{33}$$

Alternately,

$$D(P_{(\pi_0,\pi_0)}(X,Y,Z)\|P_{(\pi,\pi)}(X,Y,Z)) = I(X;Y,Z). \tag{34}$$

From (30), it is evident that the binary hypothesis test in the second scenario is harder than the first. That is, the task of identifying if sequences are scrambled is easier if they are scrambled by different transformations than by the same.

Again, as the sequence of transformations are deterministic, the source i.i.d. and the channel memoryless, the result follows. The same arguments hold for the MM decoder as well. ■

Lemma 8 implies that to study the error exponent of the multi-image registration problem, it suffices to study the error exponents of the binary hypothesis tests of the form $\psi(\bar{\pi}_0, \bar{\pi})$, for all $\pi \in \Pi$.

Now we analyze the error exponents of $\psi(\bar{\pi}_0, \bar{\pi})$ to determine the error exponent of three-image registration.

Theorem 6: Let $\bar{\pi} \in \Pi^2$. Then,

$$\lim_{n\to\infty} \frac{P_{\bar{\pi}_0,\bar{\pi}}(\Phi_{\text{MM}})}{P_{\bar{\pi}_0,\bar{\pi}}(\Phi_{\text{ML}})} = 1. \tag{35}$$

Proof: Let $\bar{\pi} = (\pi, \pi)$. Using analysis similar to that of Theorem 3, we have

$$\begin{aligned} P_{\bar{\pi}_0,\bar{\pi}}(\Phi_{\text{MM}}) &= \sum_q \prod_{a\in\mathcal{X}} \prod_{b\in\mathcal{Y}} \prod_{c\in\mathcal{Z}} \left(\mathbb{P}\left[a,b,c\right]\right)^{nq(a,b,c)} \nu_{\text{MM}}(q) \\ &= P_{\bar{\pi}_1,\bar{\pi}_2}(\Phi_{\text{ML}}) \max_q \left\{\frac{\nu_{\text{MM}}(q)}{\nu_{\text{ML}}(q)}\right\}, \end{aligned}$$

where the sum is taken over the set of all possible joint types with

$$\begin{aligned} \nu_{\text{MM}}(q) &= \sum_{\mathbf{y},\mathbf{z}} \sum |T^n_{X|Y_{\pi_0},Z_{\pi_0},Y_\pi,Z_\pi}(\mathbf{y},\mathbf{z})| \\ &= \sum |T^n_{Y_{\pi_0},Z_{\pi_0},Y_\pi,Z_\pi}| \sum |T^n_{X|Y_{\pi_0},Z_{\pi_0},Y_\pi,Z_\pi}|, \end{aligned}$$

where the first sum is taken over the set of all types $T^n_{Y_{\pi_0},Z_{\pi_0},Y_\pi,Z_\pi} \subseteq T^n_{Y,Z}$ and the second sum is over the set of all conditional types $T^n_{X|Y_{\pi_0},Z_{\pi_0},Y_\pi,Z_\pi} \subseteq T^n_{X|Y,Z}$ such that Φ_{MM} gives a decision error. Similarly, $\nu_{\text{ML}}(q)$ is defined analogously for ML decision errors.

The result follows from Lemma 9. ∎

Lemma 9: There exists a sequence $\{\epsilon_n\}$ with $\lim_{n\to\infty}\epsilon_n = 0$, such that

$$\max_q \left\{ \frac{\nu_{\text{MM}}(q)}{\nu_{\text{ML}}(q)} \right\} \leq 2^{n\epsilon_n}.$$

Proof: We again observe that maximizing multiinformation is the same as minimizing joint entropy.

$$\begin{aligned}
&\max_{(\pi_1,\pi_2)\in\Pi^2} \hat{I}(X;Y_{\pi_1};Z_{\pi_2}) \\
&\quad = \max_{(\pi_1,\pi_2)} \hat{H}(X) + \hat{H}(Y_{\pi_1}) + \hat{H}(Z_{\pi_2}) - \hat{H}(X,Y_{\pi_1},Z_{\pi_2}) \\
&\quad = \hat{H}(X) + \hat{H}(\tilde{Y}) + \hat{H}(\tilde{Z}) - \min_{(\pi_1,\pi_2)\in\Pi^2} \hat{H}(X,Y_{\pi_1},Z_{\pi_2}).
\end{aligned}$$

Thus, for $V = (Y,Z)$, solving $\psi(\bar{\pi}_0,\bar{\pi})$ using MM is the same as minimizing the joint empirical entropy $\hat{H}(X,V_{\hat{\pi}}), \hat{\pi}\in\{\pi_0,\pi\}$. This is the same as the problem in Lemma 5 with a larger output alphabet. Thus, the result follows. ∎

Theorem 7:

$$\mathcal{E}(\Phi_{\text{MM}}) = \mathcal{E}(\Phi_{\text{ML}}). \tag{36}$$

Proof: The result follows directly from Lemmas 7 and 8, and Theorem 6. ∎

The above results can directly be extended to the case of registering any finite number of copies of an image. This indicates that the max multiinformation based multi-image registration is asymptotically optimal and universal.

V. Joint Clustering and Registration

We now extend our focus to the original problem of joint clustering and registration of multiple images. As mentioned in the image model, we know that images corresponding to different scenes are independent. We are now interested in estimating the partition and registering images within each cluster.

The ML estimates are obtained as

$$\left(\hat{P}_{\text{ML}},\hat{\pi}_{\text{ML}}\right) = \arg\max_{P\in\mathcal{P},\pi\in\Pi^m} \prod_{C\in P}\prod_{i=1}^{n} \mathbb{P}\left[\mathbf{X}^{(C)}_{\pi^{(C)}(i)}\right], \tag{37}$$

where the probability is computed by averaging over scene configurations and corresponding channel model.

However, in this case it is worth noting that the ML estimates are not Bayes optimal as the prior on the set of possible permutations is not uniform. For instance, if $p_i = \frac{1}{\ell}$ for all $i\in[\ell]$, then for any partition $P\in\mathcal{P}$,

$$\mathbb{P}\left[P^* = P\right] = \frac{m!}{k!}\frac{1}{\ell^m},$$

where k is the number of clusters in the partition P. Hence, in this context, the Bayes optimal test is the likelihood ratio test, i.e., the maximum a posteriori (MAP) estimate Φ_{MAP} given by

$$\left(\hat{P}_{\text{MAP}},\hat{\pi}_{\text{MAP}}\right) = \arg\max_{P\in\mathcal{P},\pi\in\Pi^m} \mathbb{P}\left[P^* = P\right] \prod_{C\in P}\prod_{i=1}^{n} \mathbb{P}\left[\mathbf{X}^{(C)}_{\pi^{(C)}(i)}\right], \tag{38}$$

Additionally, both ML and MAP estimates require the knowledge of channel and prior distributions, and are also hard to compute. Hence we focus on the design of computationally efficient and exponentially consistent algorithms.

A. Clustering Criterion

Designing any unsupervised clustering algorithm requires a criterion that quantifies similarity. To this end, algorithms are designed in contexts such as knowing the number of clusters, or quantifying a notion of closeness for objects of the same cluster.

Here, we know that the similarity criterion is dependence of pixel values among images of the same cluster. To deal with this, one could adopt any of the following methods.

(B1) *ϵ-likeness:* A given source and channel model for images, $\mathbf{X}^{[m]}$, is said to satisfy *ϵ-likeness* criterion if

$$\min_{P^* \in \mathcal{P}} \min_{P \not\preceq P^*, P \not\succeq P^*} I_C^{P^*}(X^{[m]}) - I_C^P(X^{[m]}) \geq \epsilon.$$

(B2) *Given number of clusters:* Given the number of clusters k in the set of images, we can define an exponentially consistent universal clustering algorithm.

(B3) *Non-exponentially consistent:* Any two images $\mathbf{X}, \mathbf{Y}$ that belong to different clusters are independent, i.e., $I(X;Y) = 0$. Hence, using a threshold γ_n, decreasing with n, we can define a consistent clustering algorithm which however lacks exponential consistency such as in [5].

(B4) *Hierarchical clustering:* If it suffices to create the tree of similarity among images through hierarchical clustering, then we can define a consistent algorithm that determines such topological relation among images.

Criterion (B1) imposes a restriction on the allowed priors and channels W and may be interpreted as the capacity. On the other hand (B2) restricts the space of partitions much like k-means clustering. Criterion (B3) focuses on the design of asymptotically consistent universal clustering algorithms albeit with sub-exponential consistency. Finally, criterion (B4) aims to develop a topology of independence-based similarity relations among images. We address the clustering problem from each of these criteria.

B. Multivariate Information Functionals for Clustering

Clustering random variables using information functionals has been well studied [2], [3]. Since clustering is independence-based, we adopt the minimum partition information (MPI) framework.

Definition 13 (Partition Information): For a set of random variables $\{Z_1, \ldots, Z_n\}$ with $n > 1$, and a partition P of $[n]$ with $|P| > 1$, the partition information is defined as

$$I_P\left(Z^{[n]}\right) = \frac{1}{|P| - 1}\left[\left(\sum_{C \in P} H\left(Z^C\right)\right) - H\left(Z^{[n]}\right)\right]. \tag{39}$$

The chain rule for partition information is given as

$$I_P(Z^{[n]}) = \frac{1}{k-1}\sum_{i=1}^{k-1} I\left(Z^{C_i}; Z^{\{\cup_{j=i+1}^{k} C_j\}}\right), \tag{40}$$

where $P = \{C_1, \ldots, C_k\}$. In particular, for $P = \{\{1\}, \ldots, \{n\}\}$, i.e., each object constitutes its own cluster,

$$I_P(Z^{[n]}) = \frac{1}{n-1} I_M(Z^{[n]}).$$

The minimum residual independence in a collection of random variables is quantified by the minimum partition information.

Definition 14 (Minimum Partition Information): For a set of random variables $Z^{[n]}$ with $n > 1$, the minimum partition information is

$$I_{\text{MPI}}(Z^{[n]}) = \min_{P \in \mathcal{P}'} I_P(Z^{[n]}), \tag{41}$$

where $\mathcal{P}'$ is the set of all partitions of $[n]$ such that for all $P \in \mathcal{P}'$, $|P| > 1$.

We note here that the partition that minimizes the partition information is not necessarily unique.

Definition 15 (Fundamental Partition): Let

$$\tilde{\mathcal{P}} = \arg\min_{P \in \mathcal{P}'} I_P\left(Z^{[n]}\right).$$

Then, the *fundamental partition* of the set of random variables $Z^{[n]}$ is the finest partition $P^* \in \tilde{\mathcal{P}}$. That is, $P^* \in \tilde{\mathcal{P}}$ such that for all $P \in \tilde{\mathcal{P}}$, $P^* \preceq P$.

For the image registration problem at hand, let $X^{(i)}$ be the random variable representing a pixel of image i. Then, the images satisfy the following properties.

Lemma 10: If the images are self-aligned, i.e., $\bar{\pi}^* = \bar{\pi}_0$, and there exists at least two clusters, then the correct clustering is given by the fundamental partition and the minimum partition information is zero.

Proof: Let P^* be the correct clustering. First we note that

$$I_{P^*}(X^{[m]}) = 0,$$

as images corresponding to different clusters are independent. Further we know that the partition information is non-negative. Hence,

$$P^* \in \arg\min_{P \in \mathcal{P}'} I_P(X^{[m]}).$$

Let $\tilde{\mathcal{P}} = \{P' \in \mathcal{P}' : P' \succeq P^*\}$ and let $P \in \mathcal{P}' \backslash \tilde{\mathcal{P}}$. Then there exist clusters $C_1, C_2 \in P$ such that there exist images in each cluster corresponding to the same image. This indicates that $I(X^{C_1}; X^{C_2}) > 0$. Thus, from (40),

$$I_{P^*}(X^{[m]}) \geq I(X^{C_1}; X^{C_2}) > 0.$$

This indicates that the set of all partitions that minimize the partition information is the set $\tilde{P}$. Hence the fundamental partition is the correct clustering. ∎

Lemma 11: Let $\bar{\pi} = (\pi_1, \ldots, \pi_m)$ be the estimated transformations of the images and let $\bar{\pi}^*$ be the correct registration. Then, if $\bar{\pi} \neq \bar{\pi}^*$, then the fundamental partition $\hat{P}$ of $\{X_{\pi_1}^{(1)}, \ldots, X_{\pi_m}^{(m)}\}$ satisfies $\hat{P} \prec P^*$, where P^* is the correct clustering.

Proof: We first note that images that correspond to different scenes are independent of each other, irrespective of the transformation. Second, an image that is incorrectly registered appears independent of any other image corresponding to the same scene. This in turn yields the result. ∎

These properties provide a natural estimator for joint clustering and registration, provided the information values can be computed accurately.

Corollary 1: Let $\hat{P}_{\bar{\pi}}$ be the fundamental partition corresponding to the estimated transformation vector $\bar{\pi}$. Then $(P^*, \bar{\pi}^*)$ is the densest partition in $\{\hat{P}_{\bar{\pi}} : \bar{\pi} \in \Pi^m\}$ and the corresponding transformation vector.

Proof: This follows directly from Lemmas 10 and 11. ∎

It is worth noting here that the partition information is submodular and constitutes a Dilworth truncation lattice through the residual entropy function [3]. Thus, the fundamental partition may be obtained efficiently using submodular function minimization in polynomial time [25]. We do not get into the implementation aspects of the algorithm in this work as we focus on the consistency and the error exponents of algorithms.

We now simplify the clustering criterion for better understanding and ease of analysis. Let us define a multivariate information called *cluster information* to quantify intra-cluster information.

Definition 16 (Cluster Information): The *cluster information* of $Z^{[n]}$ for partition $P = \{C_1, \ldots, C_k\}$ of $[n]$ is

$$I_C^{(P)}(Z_1; \ldots; Z_n) = \sum_{C \in P} I_M(Z^C). \tag{42}$$

Lemma 12: For any $P \in \mathcal{P}'$, and random variables $X^{[m]}$,

$$I_P(X^{[m]}) = \frac{\left(I_M(X_1; \ldots; X_m) - I_C^{(P)}(X_1; \ldots; X_m)\right)}{(|P| - 1)}. \tag{43}$$

Proof: The result follows directly from the definitions. ∎

Corollary 2: The fundamental partition maximizes the cluster information if the minimum partition information is 0.

Proof: Follows from Lemma 12 and non-negativity of partition information. ∎

As we seek to identify the partition that maximizes the cluster information, it is worthwhile understanding its computational complexity.

Lemma 13: The clustering information of a set of random variables $\{Z_1, \ldots, Z_n\}$ is supermodular.

Algorithm 1 ϵ-like Clustering, $\Phi_{\mathrm{C}}\left(\mathbf{X}^{(1)},\ldots,\mathbf{X}^{(m)},\epsilon\right)$

for all $\bar{\pi}\in\Pi^m$ **do**
  Compute empirical pmf $\hat{p}\left(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)}\right)$
  $\tilde{I}=\max_{P\in\mathcal{P}}\hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)})$
  $P_{\pi}=\text{Finest}\left\{P:\hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)})\geq\tilde{I}-\frac{\epsilon}{2}\right\}$
end for
$(\hat{P},\hat{\pi})=\left\{(P,\bar{\pi}):P=\hat{P}_{\bar{\pi}}\succeq\hat{P}_{\bar{\pi}'},\ \text{for all }\bar{\pi}'\neq\bar{\pi}\right\}$

Proof: The clustering information may be decomposed as

$$I_C^{(P)}(Z_1;\ldots;Z_n)=\sum_{i=1}^{n}H(Z_i)-\sum_{C\in P}H(Z^C).$$

It is well established that the entropy function is submodular. since the clustering information is the difference between a modular function and a submodular function, it is supermodular. ∎

Supermodular function maximization can be performed efficiently in polynomial time. From Lemma 13 we can see that given the joint distribution, the fundamental partition may be obtained efficiently.

We now use these observations to define clustering algorithms using plug-in estimates, under each clustering criterion.

C. Consistency of Plug-in Estimates

Since the multivariate information estimates are now obtained empirically, we first show that the plug-in estimates are exponentially consistent.

Lemma 14: The plug-in estimates of partition and cluster information are exponentially consistent.

Proof: From the chain rule (40) and the union bound, we know that for all $\epsilon>0$, there exists $\delta_\epsilon>0$ and constant c such that

$$\mathbb{P}\left[\left|\hat{I}_P\left(X^{[m]}\right)-I_P\left(X^{[m]}\right)\right|>\epsilon\right]\leq c\exp\left(-n\delta_\epsilon\right), \tag{44}$$

from Theorem 1.

The result for the cluster information estimate follows from (42), the union bound, and the exponential consistency of plug-in estimates of multiinformation as

$$\mathbb{P}\left[\left|\hat{I}_C^{(P)}\left(X^{[m]}\right)-I_C^{(P)}\left(X^{[m]}\right)\right|>\epsilon\right]\leq C\exp\left(-n\delta_\epsilon\right), \tag{45}$$

where $C=2m$ and $\delta_\epsilon=\frac{\epsilon^4}{32m^4|\mathcal{X}|^{2m}\log 2}+o(1)=\theta\epsilon^4+o(1)$. ∎

We now use the plug-in estimates to define universal clustering algorithms.

D. ϵ-like clustering

Using Corollary 2 we define ϵ-like clustering algorithm, Algorithm 1. Here Finest $\{\cdot\}$ refers to the finest partition in the set.

Lemma 15: Let the source and channel be ϵ-alike, for some $\epsilon>0$. Then, Φ_C is exponentially consistent for joint clustering and registration.

Proof: First let the estimated transformation be correct, i.e., $\bar{\pi}=\pi^*$. Let $\tilde{P}=\arg\max_{P\in\mathcal{P}}\hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)})$. Then, there exists constants c,δ_ϵ such that

$$\begin{aligned}\mathbb{P}\left[\hat{I}_C^{(P^*)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_1}^{(m)})\leq\tilde{I}-\tfrac{\epsilon}{2}\right]&\leq\mathbb{P}\left[|\tilde{I}-\hat{I}_C^{(P^*)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)})|\geq\tfrac{\epsilon}{2}\right]\\&\leq\mathbb{P}\left[|\hat{I}_C^{(\tilde{P})}(X_{\bar{\pi}}^{[[m]]})-I_C^{(\tilde{P})}(X_{\bar{\pi}}^{[[m]]})|+|\hat{I}_C^{(P^*)}(X_{\bar{\pi}}^{[[m]]})-I_C^{(P^*)}(X_{\bar{\pi}}^{[[m]]})|\geq\tfrac{\epsilon}{2}\right] &(46)\\&\leq 2c\exp\left(-n\delta_{\epsilon/4}\right), &(47)\end{aligned}$$

where (46) follows from the triangle inequality and the fact that the multiinformation is maximized by the correct clustering, and (47) follows from the union bound and Lemma 14.

Further, for any $P \in \mathcal{P}$ such that $P \not\succeq P^*$ and $P \not\preceq P^*$, we have

$$\mathbb{P}\left[\hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_1}^{(m)}) \geq \tilde{I} - \tfrac{\epsilon}{2}\right] \leq \mathbb{P}\left[\hat{I}_C^{(P^*)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)}) - \hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)}) \leq \tfrac{\epsilon}{2}\right] \tag{48}$$

$$\leq \mathbb{P}\left[\left|\left(\hat{I}_C^{(P^*)}(X_{\bar{\pi}}^{[[m]]}) - I_C^{(P^*)}(X_{\bar{\pi}}^{[[m]]})\right) - \left(\hat{I}_C^{(P)}(X_{\bar{\pi}}^{[[m]]}) - I_C^{(P)}(X_{\bar{\pi}}^{[[m]]})\right)\right| \geq \tfrac{\epsilon}{2}\right] \tag{49}$$

$$\leq 2c\exp\left(-n\delta_{\epsilon/4}\right), \tag{50}$$

where (48) follows from the fact that $\tilde{I}$ is the maximum empirical cluster information and (49) follows from the ϵ-likeness criterion. Finally, (50) follows from the triangle inequality, union bound and Lemma 14.

From (50), (47), and the union bound we know that

$$\mathbb{P}\left[\hat{P}_{\bar{\pi}} \neq P^*\right] \leq 4c\exp\left(-n\delta_{\epsilon/4}\right). \tag{51}$$

Now, invoking Lemma 11, we know from similar analysis as above, that the densest fundamental partition is exactly P^*. More specifically, for any $\bar{\pi} \neq \bar{\pi}^*$, the equivalent fundamental partition is finer than P^*. This in turn indicates that

$$P_e(\Phi_C) \leq 4c|\Pi|\exp\left(-n\delta_{\epsilon/4}\right), \tag{52}$$

owing to the union bound, and (51). This proves the result. ∎

From Lemma 12, we can additionally see that an equivalent clustering algorithm can be defined in terms of the MPI functional as well and that it is also exponentially consistent, provided the underlying partition has at least two clusters.

E. K-info clustering

Alternatively under (B2), i.e., given number of clusters K in the set of images, let $\mathcal{P}_K \subset \mathcal{P}$ be the set of all partitions consisting K clusters. Then, much in the spirit of K-means clustering, we define the K-info clustering estimate as

$$(\hat{P}, \hat{\pi}) = \arg\max_{P\in\mathcal{P}_K, \bar{\pi}\in\Pi^m} \hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_1}^{(m)}) \tag{53}$$

Again, this can be extended directly to use the MPI functional.

Lemma 16: Given the number of clusters K in the set, the K-info clustering estimates are exponentially consistent.

Proof: Let $P_{\bar{\pi}} = \arg\max_{P\in\mathcal{P}_K} I_C^{(P)}(X_{\bar{\pi}}^{[m]})$ and $\hat{P}_{\bar{\pi}} = \arg\max_{P\in\mathcal{P}_K} \hat{I}_C^{(P)}(X_{\bar{\pi}}^{[m]})$. Then, for any $\bar{\pi} \in \Pi^m$,

$$\mathbb{P}\left[\hat{P}_{\bar{\pi}} \neq P_{\bar{\pi}}\right] = \mathbb{P}\left[\hat{I}_C^{(\hat{P}_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]}) > I_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]})\right]$$

$$\leq 2(|\mathcal{P}_K| - 1)c\exp\left(-n\delta_{\epsilon_{\bar{\pi}}}\right), \tag{54}$$

where (54) follows from the union bound and Lemma 14, for $\epsilon_{\bar{\pi}} = I_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]}) - \max_{P\neq P_{\bar{\pi}}} I_C^{(P)}(X_{\bar{\pi}}^{[m]})$.

Next, for any $\bar{\pi} \neq \bar{\pi}^*$, we have

$$\mathbb{P}\left[\hat{I}_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]}) > \hat{I}_C^{(P_{\bar{\pi}^*})}(X_{\bar{\pi}^*}^{[m]})\right]$$

$$\leq \mathbb{P}\left[\left|\left(\hat{I}_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]}) - I_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]})\right) - \left(\hat{I}_C^{(P_{\bar{\pi}^*})}(X_{\bar{\pi}^*}^{[m]}) - I_C^{(P_{\bar{\pi}^*})}(X_{\bar{\pi}^*}^{[m]})\right)\right| \geq I_C^{(P_{\bar{\pi}^*})}(X_{\bar{\pi}^*}^{[m]}) - I_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]})\right]$$

$$\leq 2c\exp\left(-n\delta_{\tilde{\epsilon}_{\bar{\pi}}}\right), \tag{55}$$

where $\tilde{\epsilon}_{\bar{\pi}} = I_C^{(P_{\bar{\pi}^*})}(X_{\bar{\pi}^*}^{[m]}) - I_C^{(P_{\bar{\pi}})}(X_{\bar{\pi}}^{[m]})$.

Finally, the result follows from (54), (55), the union bound, and the fact that the number of possible transformations is polynomial in n. ∎

Algorithm 2 Thresholded Clustering, $\Phi_{\mathrm{T}}\left(\mathbf{X}^{(1)},\ldots,\mathbf{X}^{(m)},\alpha\right)$

$\gamma_n \leftarrow c_1 n^{-\alpha}$, for some constant $c_1 > 0$
for all $\bar{\pi} \in \Pi^m$ **do**
  Compute empirical pmf $\hat{p}\left(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)}\right)$
  $\tilde{I} = \max_{P\in\mathcal{P}} \hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)})$
  $P_\pi = \text{Finest}\left\{P : \hat{I}_C^{(P)}(X_{\pi_1}^{(1)},\ldots,X_{\pi_m}^{(m)}) \geq \tilde{I} - \gamma_n\right\}$
end for
$(\hat{P},\hat{\pi}) = \left\{(P,\bar{\pi}) : P = \hat{P}_{\bar{\pi}} \succeq \hat{P}_{\bar{\pi}'},\ \text{for all } \bar{\pi}' \neq \bar{\pi}\right\}$

Algorithm 3 Hierarchical Clustering, $\Phi_{\mathrm{H}}\left(\mathbf{X}^{(1)},\ldots,\mathbf{X}^{(m)}\right)$

$\hat{\pi} = \arg\max_{\bar{\pi}\in\Pi^m} \hat{I}_M(X_{\bar{\pi}}^{[m]})$
$\mathbf{Y}^{[m]} \leftarrow X_{\hat{\pi}}^{[m]}$
$P \leftarrow \{[m]\}$, $\hat{P}(1) \leftarrow P$
for $k = 2$ **to** m **do**
  for all $C \in P$ **do**
    $J_C \leftarrow \max_{\tilde{P}\in\mathcal{P}_2(C)} I_C^{(\tilde{P})}(Y^C)$
    $\tilde{P}_C \leftarrow \arg\max_{\tilde{P}\in\mathcal{P}_2(C)} I_C^{(\tilde{P})}(Y^C)$
  end for
  $\tilde{C} \leftarrow \arg\max_{C\in P} J_C$
  $P \leftarrow P \backslash \tilde{C}$
  $P \leftarrow P \cup \tilde{P}_{\tilde{C}}$
  $\hat{P}(k) \leftarrow P$
end for

F. Clustering with sub-exponential consistency

We now address clustering under criterion (B3). Specifically, we focus on the design of a universal clustering algorithm that is not aware of the underlying parameters of ϵ-likeness and number of clusters K.

We know that independence characterizes dissimilarity of images. Given a source and channel model, there exists $\epsilon > 0$ such that the set of images are ϵ-alike. However, the value of ϵ is unknown to the decoder. Since the plug-in estimates are exponentially consistent, we adapt Algorithm 1 to work with a dynamic threshold that takes into consideration the resolution (the number of pixels) of the image. The threshold-based clustering algorithm is given in Algorithm 2.

Lemma 17: Let $\alpha \in (0, \frac{1}{4})$. Then, the thresholded clustering algorithm, Φ_T is asymptotically consistent.

Proof: Let the given source and channel model be ϵ-alike. Since γ_n is decreasing with n, there exists $N_\epsilon < \infty$ such that for all $n > N_\epsilon$, $\gamma_n < \epsilon$. The proof now follows analogous to the proof of Lemma 15, from the observation that $n\delta_{\gamma_n} \to \infty$ as $n \to \infty$. ∎

Hence can see that, at the expense of exponential consistency, we can design a reliable universal clustering scheme that only uses the independence for clustering and registering the images.

G. Hierarchical Clustering

Finally we consider clustering according to criterion (B4). Here, we aim to establish the hierarchical clustering relation among images to establish the phylogenetic tree for the images.

From the nature of independence among dissimilar images, we observe a natural algorithm to construct such hierarchical clustering as shown in Algorithm 3.

Fundamentally, at each stage, the algorithm splits into two, one cluster from the existing partition. More specifically, the cluster that has the most impact upon such a split in terms of the cluster information is chosen at each stage. We observe a natural notion of association dependent on the degree of intra-cluster dependence.

Lemma 18: Consider a set of images with K clusters. Then, the resulting partition at iteration $k = K$ of Φ_H is almost surely the correct clustering.

Alternately, if the system is ϵ-alike, then for sufficiently large n, at iteration $k = |P^*|$ the estimated partition is correct with arbitrarily high probability.

Further, the estimate of registration is exponentially consistent.

Proof: The proof follows directly from the exponential consistency of the other clustering and registration algorithms. ■

H. Computational and Sample Complexity

Although not the principal focus of this paper, we make a brief note of the computational and sample complexity of the algorithms defined here.

We have established that the fundamental partitions for cluster information may be obtained efficiently for any chosen transformation vector, given the joint distribution. However, exploring the neighborhood of the maximizing partition is a harder problem to address and it is not clear if there exists an efficient algorithm, other than exhaustive search, to do this step.

Further, identifying the correct transformation still involves an exhaustive search in the absence of additional information on the nature of the transformations, dependency in the transformations across different images, or underlying correlations across pixels of an image. More specifically, since $|\Pi| = O(n^{\alpha})$, the computational complexity of the exhaustive search is $O(n^{\alpha m})$, i.e., exponential in the number of images.

We note that for a finite number of images to be clustered and registered, this method of exploration is typically sufficient. However, the question of improving computational complexity at the expense of universality, and possibly accuracy, is certainly interesting, and worth further study.

With regard to sample complexity, as we deal with the computation of information measures across a set of m random variables, to suppress the bias in the plug-in estimates, the sample complexity of the clustering and registration algorithms is $O(r^m)$. This translates to the requirement of high resolution images. Additional information on structure in the underlying image model could be used to reduce the sample complexity.

VI. Conclusion

In this work, we study the problem of joint clustering and registration of images in a universal setting. We perform an information-theoretic analysis of the MMI method for registering images and define appropriate multivariate information functionals to perform efficient and asymptotically optimal clustering and registration of images.

The information-theoretic framework established here to study the performance of the algorithm may potentially be extended to more generic image models that incorporate inter-pixel correlations. Specific knowledge about image transformations could further help improve the performance of the defined algorithms.

Acknowledgment

The authors would like to thank Vinith Misra for generic insight into universal information processing and specific intuitive outlook into the image registration problem presented in [26].

References

[1] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," *J. Mach. Learn. Res.*, vol. 6, pp. 1705–1749, Oct. 2005.

[2] K. Nagano, Y. Kawahara, and S. Iwata, "Minimum average cost clustering," in *Advances in Neural Information Processing Systems 23*, J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, Eds. MIT Press, 2010, pp. 1759–1767.

[3] C. Chan, A. Al-Bashabsheh, J. B. Ebrahimi, T. Kaced, and T. Liu, "Multivariate mutual information inspired by secret-key agreement," *Proc. IEEE*, vol. 103, no. 10, pp. 1883–1913, Oct. 2015.

[4] V. Misra and T. Weissman, "Unsupervised learning and universal communication," in *Proc. 2013 IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 261–265.

[5] R. K. Raman and L. R. Varshney, "Universal clustering via crowdsourcing," arXiv:1610.02276, Oct. 2016.

[6] M. E. Leventon and W. E. L. Grimson, "Multi-modal volume registration using joint intensity distributions," in *Int. Conf. on Medical Image Comp. and Comp.-Assisted Intervention*, 1998, pp. 1057–1066.

[7] H.-M. Chan, A. C. S. Chung, S. C. H. Yu, A. Norbash, and W. M. Wells, "Multi-modal image registration by minimizing Kullback-Leibler distance between expected and observed joint class histograms," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognition (CVPR'03)*, vol. 2, Jun. 2003, pp. II–570.

[8] W. K. Pratt, "Correlation techniques of image registration," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-10, no. 3, pp. 353–358, May 1974.

[9] P. Viola and W. M. Wells III, "Alignment by maximization of mutual information," *Int. J. Comput. Vis.*, vol. 24, no. 2, pp. 137–154, 1997.

[10] A. Goshtasby, G. C. Stockman, and C. V. Page, "A region-based approach to digital image registration with subpixel accuracy," *IEEE Trans. Geosci. Remote Sens.*, no. 3, pp. 390–399, 1986.

[11] H. S. Alhichri and M. Kamel, "Image registration using the Hausdorff fraction and virtual circles," in *Proc. IEEE Int. Conf. Image Process. (ICIP 2001)*, vol. 2, Oct. 2001, pp. 367–370 vol.2.

[12] B. Zitova and J. Flusser, "Image registration methods: a survey," *Image Vision Comput.*, vol. 21, no. 11, pp. 977–1000, 2003.

[13] M. Xu, H. Chen, and P. K. Varshney, "Ziv-Zakai bounds on image registration," *IEEE Trans. Signal Process.*, vol. 57, no. 5, pp. 1745–1755, May 2009.

[14] C. Aguerrebere, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5707–5722, Nov. 2016.

[15] V. D. Goppa, "Nonprobabilistic mutual information without memory," *Probl. Control Inf. Theory*, vol. 4, no. 2, pp. 97–102, 1975.

[16] H. D. Tagare and M. Rao, "Why does mutual-information work for image registration? a deterministic explanation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 6, pp. 1286–1296, Jun. 2015.

[17] L. Zollei and W. M. Wells, "On the optimality of mutual information as an image registration objective function," in *Proc. IEEE Int. Conf. Image Process. (ICIP 2009)*, Nov. 2009, pp. 189–192.

[18] J. Stein, J. Ziv, and N. Merhav, "Universal delay estimation for discrete channels," *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 2085–2093, Nov. 1996.

[19] P. Ngo, Y. Kenmochi, N. Passat, and H. Talbot, "On 2D constrained discrete rigid transformations," *Ann. Math. Artif. Intell.*, vol. 75, no. 1, pp. 163–193, 2015.

[20] L. A. Shepp and S. P. Lloyd, "Ordered cycle lengths in a random permutation," *Trans. Am. Math. Soc.*, vol. 121, no. 2, pp. 340–357, 1966.

[21] I. Csiszár, "The method of types," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2505–2523, Oct. 1998.

[22] P. Whittle, "Some distribution and moment formulae for the Markov chain," *J. R. Stat. Soc. Ser. B. Methodol.*, vol. 17, no. 2, pp. 235–242, 1955.

[23] P. Netrapalli, S. Banerjee, S. Sanghavi, and S. Shakkottai, "Greedy learning of markov network structure," in *Proc. 48th Annu. Allerton Conf. Commun. Control Comput.*, Sept 2010, pp. 1295–1302.

[24] M. Studený and J. Vejnarová, "The multiinformation function as a tool for measuring stochastic dependence," in *Learning in Graphical Models*, M. I. Jordan, Ed. Dodrecht: Kluwer Academic Publishers, 1998, pp. 261–297.

[25] C. Chan, A. Al-Bashabsheh, Q. Zhou, T. Kaced, and T. Liu, "Info-clustering: A mathematical theory for data clustering," *IEEE Trans. Mol. Biol. Multi-Scale Commun.*, vol. 2, no. 1, pp. 64–91, Jun. 2016.

[26] V. Misra, "Universal communication and clustering," Ph.D. dissertation, Stanford University, Jun. 2014.